arXiv:physics/0110067v1 [physics.bio-ph] 23 Oct 2001

Hamming distance geometry of a protein conformational space.

Application to the clustering of a 4 ns molecular dynamics trajectory of the HIV-1 integrase catalytic core.

Cyril Laboulais[2], Mohammed Ouali[3], Marc Le Bret[2]
and Jacques Gabarro-Arpa[1 2 *]

[1)] Ecole Normale Supérieure / C.N.R.S.
Laboratoire Interdisciplinaire de Géometrie Appliquée
45, rue d'Ulm, 75230 Paris cedex (France)

[2)] LBPA, C.N.R.S. UMR 8532, Ecole Normale Supérieure de Cachan
61, Avenue du Président Wilson, 94235 Cachan cedex (France)

[3)] University of Wales, Department of Computer Science
Group of Computational Biology, Penglais, Aberystwyth SY23 3DB (Wales)

Abstract

Protein structures can be encoded into binary sequences[1], these are used to define a Hamming distance in conformational space: the distance between two different molecular conformations is the number of different bits in their sequences.

Each bit in the sequence arises from a partition of conformational space in two halves. Thus, the information encoded in the binary sequences is also used to characterize the regions of conformational space visited by the system.

We apply this distance and their associated geometric structures, to the clustering and analysis of conformations sampled during a 4 ns molecular dynamics simulation of the HIV-1 integrase catalytic core.

The cluster analysis of the simulation shows a division of the trajectory into two segments of 2.6 and 1.4 ns length, which are qualitatively different: the data points to the fact that equilibration is only reached at the end of the first segment.

Some length of the paper is devoted to compare the Hamming distance to the r.m.s. deviation measure. The analysis of the cases studied so far, shows that under the same conditions the two measures behave quite differently, and that the Hamming distance appears to be more robust than the r.m.s. deviation.

* Corresponding author. Electronic Address: gabarro@dmi.ens.fr

Introduction

Proteins are mesoscopic systems with a complex potential energy surface that can be characterized by its locally stable conformations, which are local potential energy surface minima (PESM). There is evidence that proteins move between these conformational substates and that they can exist in many levels[2,3]. Because of their complexity, methods for analyzing these multiple minima can be very useful.

Molecular dynamics simulations (MD) is one of the available techniques for simulating the wandering of such systems across their potential energy surface. To detect the energy substates visited by the resulting trajectory we make the assumption that the structures found in PESMs are more related to each other than to conformations in other minima; thus, we follow the approach that consists in using clustering algorithms to partition the trajectory into groups of similar conformers, as has been done by a number of authors[1,4-7]. This allows to extract useful information and at the same time reducing to tractable dimensions the ever growing number of data in these computer experiments.

To put related conformations together we need a distance in conformational space. In this paper we use a Hamming distance between conformations[1], which is built as follows:

1) we assume that there is an order relation among the N atoms in the molecule,

2) we form the set of all $P = \binom{N}{4}$ ordered 4-tuples of atoms,

3) for each ordered 4-tuple of atoms $\{a, b, c, d\}$, we calculate the signed volume of the 3-simplex determined by the position of these four atoms:

$$V = 1/6 \begin{vmatrix} x_1^a & x_2^a & x_3^a & 1 \\ x_1^b & x_2^b & x_3^b & 1 \\ x_1^c & x_2^c & x_3^c & 1 \\ x_1^d & x_2^d & x_3^d & 1 \end{vmatrix} \qquad (1)$$

4) from it we define

$$\chi(a,b,c,d) = \begin{cases} 1 & V > 0 \\ 0 & V = 0 \\ -1 & V < 0 \end{cases} \qquad (2)$$

the set of all χ is called the **chirotope**, it is a binary sequence (if we discard the 0 value as it corresponds to a set of null measure). It defines an equivalence relation among conformations: two conformations are equivalent if they have the same chirotope. When comparing two conformations, whenever a 4-tuple χ value has different signs we call it a **mutation**: the number of mutations between two conformations is a Hamming distance.

The Hamming distance and its associated geometrical structures, as a tool for peering at the energy landscape of a protein, is the subject of this paper. The Hamming distance has a number of interesting features which can be summarized as follows:

- It is a true distance obeying the triangle inequality.

- It is sensitive to chirality.

- It is computed over a great number of binary topological descriptors.

- 4-tuples give meaningful geometric objects in the space of conformations, since each one determines a partition of this space into positive and negative halves.

The latter can help to put limits on the regions occupied by a particular trajectory segment, or to determine the crossing frequency for some boundaries, as we shall see below.

In what follows we develop our subject around the study of a 4 ns MD trajectory of the HIV-1 integrase catalytic core.

Molecular dynamics

The simulated molecule is the HIV-1 integrase catalytic core (residues 50-212 of the integrase); it contains the single mutation F185H and is studied in the presence of its cofactor, an Mg^{++} ion. It was taken from the molecule C of the 1bl3 PDB file[8] for which the coordinates of all the residues in the catalytic loop are available. The end residues (50, 211 and 212) were added to the structure which was then minimized using our quasi-newtonian minimizer[9].

As the total charge of the protein with the Mg^{++} cation is +1, to ensure electric neutrality an extra Cl^- anion was added into the simulation box. The 13 nearest water molecules to the Mg^{++} cation were positioned as in the PDB file. The EDIT module of AMBER[10] was modified so that a box of given size $50 \times 60 \times 50$ $\mathring{A}^3$ could be filled with TIP3P water molecules having a density of 1, which amounted to a total of 3822+13 water molecules for 2517+2 solute atoms.

The SANDER module of AMBER was used to simulate the dynamics. The force field[11] was complemented with the parameters of Mg^{++} and Cl^- ions[12]. Long range coulombic interactions were calculated using the particle mesh Ewald method[13], with a 10 $\mathring{A}$ cut off, and grid size $50 \times 60 \times 50$.

Covalent bonds containing a proton were constrained to their equilibrium length using the SHAKE algorithm[14]. The simulation was run in 2 fs steps, at 300 K constant temperature and 1 atm constant pressure. No correction was applied for the neglected long-range van der Waals interactions because they are expected to be small.

The protocol to set the water around the protein and ions, and starting the simulation was discussed in ref. 15. It consists of the following steps:

- 5 ps of water heating from 0 K to 300 K with the ions and protein constrained to their initial positions,
- next the system was restarted at 10 K and heated to 300 K, in intervals of 25 K and 1 ps. As in the previous step only water was allowed to move.
- After 3.6 ps at 300 K, the system was restarted at 10 K and heated again to 300 K in intervals of 50 K and 3 ps without any constraints.
- Finally, after reseting the velocities, we let the system evolve in an equilibration step lasting 13 ps at 300 K.

Once this protocol was finished the simulation was started and let run for a total simulated time of 4 ns, conformations were sampled every 0.1 ps giving a total of 40000 coordinate frames. The data were visualized and analysed using the program Cadira[16].

Analysis of the space of conformations

A lot of information can be encoded in the chirotope. In this paper we will use it to extract information concerning the trajectory behaviour in the regions of conformational space it crosses.

Each 4-tuple has an associated geometric structure, since for a given conformation its χ value can be either positive or negative, it determines a partition of the space of conformations into positive and negative halves. We define an **event** as the crossing by the trajectory of the boundary separating a tuple's positive from negative conformations.

For each tuple we call **major** halfspace the one in which the trajectory stays the longest and we call **minor** halfspace the other one. There is a binary **sequence** associated with each tuple of length the number of coordinate frames in the trajectory. For each frame the sequence has a value 0 or 1 depending if its conformation lies in the tuple's major or minor halfspace, respectively. An event corresponds to a change of value in the sequence.

We illustrate our point with fig. 1, where conformational space is represented, for purposes of illustration, as a two dimensional surface crossed by lines representing the separation between positive and negative halves associated with given tuples. Suppose that the trajectory evolves between two regions where it can get momentarily trapped, maybe because they correspond to PESMs, as depicted in fig. 1: the first one, **R1**, large with shallow boundary walls; the second, **R2**, somewhat smaller with a steep boundary.

The tuple separation lines within these regions fall into three classes:

- specific, those that only cross one region, lines beginning by **s** in fig. 1,

- non-specific that cross more than one region, **ns** lines,

- and those that fall in between, lines beginning by **b**.

See the legend of fig. 1 for typical example sequences associated with these tuples. We say that a given tuple is specific (non-specific) if its associated separation boundary is specific (non-specific) respectively.

Among specific tuples we can make a distinction between two extreme cases: some that cut a region across the center and others that only scratch the border, **sm** and **sw** lines in fig. 1 respectively.

Concerning the separation lines associated with **sw** tuples we see that they are only crossed for a brief lapse of time, the trajectory bouncing on the region's walls and inmediately coming back to the interior; the sequence of these tuples will only harbor a few events, with a minimum of two. A difference between regions **R1** and **R2** is that the latter has steeper walls; so, we expect that for **sw** tuples in **R2** the average number of frames lying in the tuple minor halfspace will be less than for **sw** tuples in **R1** which is shallower.

On the other hand **sm** tuples are crossed repeatedly while the trajectory stays within the region. They give rise to sequences with many mutation events within the corresponding sequence subset (see legend in fig. 1).

b tuples are crossed when the trajectory goes in between two such regions, and as can be deduced from fig. 1, they will give sequences with mostly one sign for each region.

The classes of tuples defined above are based on a purely qualitative distintion. Actually, there is no a sharp separation between them. As can be seen from fig. 3 they form a continuum.

In order to characterize a trajectory interval that stays confined in a region of the space of conformations, a previous work[1] introduced a parameter: the activity coefficient D, which can be calculated with the following procedure:

- first, we calculate the total number of events generated by the specific tuples in this cluster E_s

- second, we calculate the same quantity on every homologous cluster along the trajectory and we take the mean value $< E_s >$

- the activity coefficient D will be the ratio $E_s / < E_s >$.

The assumption is that D will exhibit a value greater than one for trajectory intervals that remain confined due to repeated bouncing on domain walls.

Clustering

The basic assumption behind the clustering methodology is that the presence of a potential well in conformational space can trap the system. So a relatively long trajectory segment can evolve in a small region of conformational space, giving a set of closely related structures that are easily detectable with an appropiate clustering algorithm[1,4-7].

In this paper we employ hierarchical data clustering, which is a process that transforms a proximity matrix into a dendogram: a tree representing a hierarchy of categories of clusters. There are other approaches to data clustering but we employ hierarchichal methods because the energy landscape of a protein is supposed to consist of an hierarchical embedding of potential wells[17-20].

Again we have two categories of hierarchical algorithms: agglomerative and divisive. Divisive algorithms were tested before[1] and they were not found to perform very well, thus in this paper we restrict ourselves to agglomerative ones. These work as follows: for a set of n elements we begin with a structure with n clusters, one per object, and grow a sequence of clusterings until all n patterns are in a single cluster. At each step two clusters are reunited if they optimize a given criteria.

We use the complete-linkage, or farthest-neighbour algorithm, in which the distance between subsets is the maximum distance between a pair of elements one in each subset. The complete-linkage hierarchy can be computed in $O(n \log^2 n)$ time[21]. One reason for choosing this particular algorithm is that it enhances cluster separation and we think it is more adequate for studying long trajectories.

One problem with clustering algorithms is that clusters will always be found even if the data are entirely random, so we need procedures that help us to determine which clusters do really correspond to PESMs. Here we employ two different methods: the first one, as explained in the previous section, the likelyhood of PESM activity for each cluster is estimated with the parameter D. Following our analysis of conformational space, a PESM should display a greater than average number of specific events specially in the range of **sw** tuples, thus a D value high above 1 can be considered a reliable indicator. In order to validate our conformational space methods, we used energy minimization as our second test procedure, as it is described in the following section.

Two distance matrices were built from the set of C_α carbons exclusively[6,7]: one based on the Hamming distance as described above, and another based on the r.m.s. deviation (RMSD). The latter is given by[22]:

$$\mathrm{RMSD} = \min_R \left(\frac{1}{N}\sum_i^N \| R.\mathrm{x}_i^a - \mathrm{x}_i^b \|^2\right)^{1/2}$$

where N is the total number of atoms, x_i^a and x_i^b denote the atomic coordinates of the ith atom of structures a and b in a common center of mass coordinate system, the distance is the minimum taken over all 3×3 rotation matrices R.

The reason for choosing the RMSD from the set of available distance measures[5], is that, as for the Hamming distance, it applies to a set of equivalence classes in conformational space: two conformations **a** and **b** are in the same equivalence class if there exists a rotation and a translation in three-dimensional space that transform **b** into **a** .

As defined above the distance between any two such equivalence classes is the minimum r.m.s. between any two conformations one in each class.

Moreover, the classes belonging to this set are obviously embedded in the equivalence classes defined by the chirotope, as lower dimensional manifolds.

Comparison with other similarity measures, like the r.m.s. difference between corresponding torsion angles, and the r.m.s. deviation of intramolecular distances[5,6], was not considered meaningful since they do not operate on cartesian conformational space, like the RMSD and the Hamming distance.

Energy Minimization

A PESM acts as an attractor on the trajectory which can get stuck in it for a while. The structures sampled along the corresponding segment of the trajectory will form a cluster of similar conformations. The contrary may not be true: a cluster of similar conformations does not necessarily indicate the presence of an energy surface minima. We use energy minimization as an alternative method to see if a cluster represents a feature of the energy lanscape.

In a first approximation, a PESM can be thought as a paraboloid shaped structure, that becomes narrower as we approach towards the minima. Thus, upon moderate minimization the distance between any two conformations in a cluster will decrease.

We performed energy minimization on structures belonging to clusters of homologous conformers using the module SANDER from AMBER[10]. Minimization was perfomed with the ewald option, up to a maximum r.m.s. gradient of 0.4, which is somewhat high. However, enhanced minimization is not of much use in our case since the effective potential surrounding the protein depends upon the organization of the solvent and ions. Continued minimization will disrupt this structure by inducing extensive reorganization of the water network.

This precludes the use of energy minimization to estimate the size of features in the energy lanscape, as well as identifying smaller minima within PESMs. But it is enough to observe a shrinking of r.m.s. deviations (RMSDS) between elements of the cluster which is an indicator of the presence of a PESM, otherwise in this paper we did not attempt at resolving the fine structure.

Results

From the 28342440 4-tuples present in our system 25 percent undergo at least one mutation during this dynamics run. Between conformations 5 ps away we observe a minimum distance of 525922 and a maximum of 2479784.

The dendogram, fig. 2, clearly separates the frames into two big clusters we call **A** and **B**: cluster **A** harbors the first 2.6 ns, cluster **B** spans 1.4 ns until the end of the run (see table I for a description of the main clusters). Cluster **A** has an activity coefficient of 6.86 and is dominated by a big subcluster: $\mathbf{A}_0$, 1.4 ns in length with a staircase structure. Interestingly $\mathbf{A}_0$ contains the clusters at the deepest level in the hierarchy which are organized around the initial structure: the cluster $\mathbf{A}_{00}$ spanning the first 45 ps has an activity coefficient of 11.07, it is embedded in cluster $\mathbf{A}_{01}$ which spans the first 255 ps with a coefficient 7.76 (see fig. 3a for a plot showing the spectrum of this cluster). This shows that the region around the initial structure exhibits PESM activity and is well separated in conformational space from the rest of the trajectory.

As we climb cluster $\mathbf{A}_0$ up to the top the activity coefficient gradually disminishes up to a minimum of 1.94 at the top, where the **sw**-tuple activity has completely disappeared from the spectrum, but there still remains a high distribution of **sm**-tuples (see fig. 3b). This trend continues past $\mathbf{A}_0$ up to **A** which also displays little **sw**-tuple activity (see fig. 3c).

It is difficult to discern a prominent feature inside $\mathbf{A}_0$; it seems as if the trajectory could not get away from the artificially created PESM around the energy minimized initial structure, and this PESM structure were slowly decaying.

Interspersed within the segment of the trajectory spanned by $\mathbf{A}_0$ there are clusters $\mathbf{A}_1$, $\mathbf{A}_3$ and $\mathbf{A}_5$ which exhibit genuine PESM activity, since they have a well defined **s**-tuple activity as well as RMSDS upon minimization. However, analysis of their tuple spectrum (fig. 3d) shows a lack of significant **b**-tuple activity for these clusters; this seems to indicate that they could arise as a bulge from the main cluster: the trajectory gets momentarily out of $\mathbf{A}_0$ and then plunges back.

Consistent with this observation is the fact that $\mathbf{A}_1$, $\mathbf{A}_3$ and $\mathbf{A}_5$ lie higher in the cluster hierarchy than $\mathbf{A}_0$, indicating that they are somewhat distant from it in conformation space.

Clusters which are relatively close to $\mathbf{A}_0$, $\mathbf{A}_2$ and $\mathbf{A}_4$ display RMSDS, but like $\mathbf{A}_0$ their **sw**-tuple activity is negligible. This trend goes up to the parent cluster $\mathbf{A}$ whose spectrum is rich in **sm**-tuples and **b**-tuples, indicating that it is well separated from the rest of the trajectory as well as some internal structure. However its **sw**-tuple is no more than average (see fig. 3c). Only the clusters around the initial energy minimized structure $\mathbf{A}_{00}$ and $\mathbf{A}_{01}$ and the bulge clusters $\mathbf{A}_1$, $\mathbf{A}_3$ and $\mathbf{A}_5$ display any **sw**-tuple activity.

This contrasts with the situation in $\mathbf{B}$ whose spectrum (depicted in fig. 3e) shows activity for every class of tuples all along the hierarchy, down to the bottom. Besides all clusters in this part of the dendogram, from $\mathbf{B}_1$ to $\mathbf{B}_5$, appear to be well separated. They represent continuous trajectory segments, and all, except $\mathbf{B}_3$, exhibit PESM activity either from the tuple spectrum or from RMSDS (see table I). $\mathbf{B}_3$ is an exception, it spans a continuous strecth of the trajectory from 2.97 to 3.53 ns. No PESM activity of any kind is detected for this cluster. Activity is only detected down the hierarchy for small clusters spanning less than 100 ps. These data are not shown since at this point our dendogram attains its limit resolution.

As before, the end of the trajectory, spanned by cluster $\mathbf{B}_5$, (see fig. 2 and fig. 3f) appears at the deepest level in cluster $\mathbf{B}$.

When we apply tuple spectrum analysis separately for clusters $\mathbf{A}$ and $\mathbf{B}$ we obtain qualitatively the same results on both, except for cluster $\mathbf{A}_5$. The activity coefficient of this cluster increases considerably, but this is not unexpected since it lies at the extremity of the trajectory segment of cluster $\mathbf{A}$. From this spectrum invariance of tuples encoding the ocupation structure of conformational space, for $\mathbf{A}$ and $\mathbf{B}$, we can conclude that there is little interference between the two regions and that the two consecutive segments of the trajectory they contain are well separated in conformational space.

In the upper right part of fig. 4 we have plotted the Hamming distance matrix, clusters $\mathbf{A}$ and $\mathbf{B}$ (see table I) are easily identified by simple visual inspection.

Comparison with r.m.s. clustering

Clustering can be performed using other measures than Hamming distance, as a comparison we have used the RMSD[22], which is a widely used similarity measure.

The dendogram which results from r.m.s. deviation clustering, using the same protocol as above, can be seen in fig. 5a. It is composed of two big clusters, superficially each one looks like $\mathbf{A}_0$ in fig. 2: they are an aggregation of small clusters of average length 100 ps, no other visual features can be discerned. A closer look to the data gives the following:

- The two main clusters correspond to discontinuous trajectory segments.
- Clusters $\mathbf{A}_{00}$ and $\mathbf{B}_5$ which contain the extremities of the trajectory, appear at branching levels relatively close to the root of the tree, moreover $\mathbf{B}_5$ comes in two separated clusters. Indeed the three is not well balanced.
- In fig. 5a we have underlined clusters which correspond to fragments of clusters identified in the dendogram of fig. 2 (with the exception of $\mathbf{A}_0$). Many previously identified clusters

(harboring continuous trajectory segments) appear fragmented, some of them scattered over many branchs.

Obviously the protocol used for the Hamming distance gives poor results when used with the RMSD.

From a previous work using the same protocol[1], it was reported more meaningful results with RMSD clustering. However, they were obtained for a 1 ns trajectory, which is four times shorter than the present one, and the protein studied was the bovine pancreatic trypsin inhibitor[23], which is about three times smaller than the one studied in the present work. Also in that occasion the authors deemed not necessary to report that the dendogram tree was poorly balanced and that the main branches harbored some discontinuities, two of the features identified above. The result seemed acceptable because the analysis of the low level clusters, using discrete geometry tests, closely matched the ones that had been previously identified using Hamming data clustering[1]. Seen in the light of the present results it seems that the work with BPTI attained the limit resolution that can be obtained with the above method, using the RMSD.

For sake of comparison, we have plotted in fig. 4 the Hamming vs the r.m.s. distance matrices. Although both display the most salient features, the Hamming distance matrix appears to be much more detailed at the fine grained level. This suggests that better results could be obtained with RMSD if resolution were improved.

One way of improving resolution is to calculate the distance matrix using more atoms than the C_α carbons, two possibilities were tested: one with all the backbone heavy atoms (carbonyl carbon, amide nitrogen and C_α), another with all the protein heavy atoms. Neither works, with the backbone heavy atoms option the resulting tree had a less monotonous appearence and had improved balancing, but again it failed to resolve the main clusters, and many smaller clusters (between 150 and 500 ps in length) corresponded to discontinuous trajectory segments.

The option involving all heavy atoms gave the worst results, the reason is that individual regions along the protein jump between conformational substates, and the combination of all these transitions distribute uniformly through conformational space, thus smearing the data[6].

It may also happen that the complete-linkage is not the best clustering algorithm available [7], at least when used together with RMSD. As in ref. 7 we have tried the average distance criterion, where the distance between two clusters is taken as the average distance between the points in one cluster and the points in the other.

The resulting dendogram, see fig. 5b, yields better results for the RMSD (with all the backbone heavy atoms option): at least the algorithm correctly identifies the main clusters **A** and **B**, so we can compare with the results from previous section. It also identifies the clusters corresponding to well defined PESMs: $\mathbf{A}_{00}$, $\mathbf{A}_{01}$, $\mathbf{A}_1$, $\mathbf{A}_5$, and every cluster on the **B** side. The **B** cluster obtained is quite close to the one in fig. 2.

The same cannot be said of the **A** subtree: to begin with it is not correctly balanced, $\mathbf{A}_{00}$ and $\mathbf{A}_{01}$ appear close to the root; also, the two halves of $\mathbf{A}_2$ and $\mathbf{A}_3$ from fig. 2, are in two separate branches; finally, $\mathbf{A}_4$ is scattered over many branches. We think that this protocol can be improved to give still better results, but in the scope of the present work, we do not see the point in getting any further.

On the other hand the average distance cluster algorithm with the Hamming distance gives a result quite similar to the one discussed in the previous section.

Discussion

After performing conformer structure clustering of an MD trajectory based on a Hamming distance, the analysis of the dendogram shows a very strong PESM at the begining of the dynamics. It surrounds the energy minimized structure created at the start of the run and lies

at the bottom of the hierarchy of a staircase shaped featureless cluster that spans 1.4 ns of the trajectory.

This PESM dominates the initial part of the simulation, analysis reveals a slow decay: as we climb the hierarchy of clusters along the stairs the activity disminishes steadily, this is due to a waning of the **sw** part of the tuple spectrum, this low **sw** activity cluster indicates an absence of sharp boundary walls. The initial PESM has progressively become shallower and structureless.

This process is ponctuated by incursions into smaller PESMs with high **sw** activity, pointing to the presence of sharp boundaries, but totally absent **b** activity (see fig. 3d), indicating that these PESMs may be a temporary bulge.

This behaviour may be attributed to the effect of the protocol used in preparing the dynamics run; the whole system is assembled from separate subsystems: the crystallographic structure for the protein and a cube containing a Monte Carlo generated distribution of water, which are relaxed separately. The result is that the final system is only locally equilibrated. The use of the Ewald summation method to calculate long range interactions contributes a great deal, since it enhances the sensitivity of the system to global conditions. The anomalous structure in the protein dendogram first half can be attributed to the fact that the system is only locally equilibrated, global equilibrium being attained much later in the dynamics run.

This kind of behaviour has already been reported[24]. There is however a difference between their results and ours: in the first half of the dynamics, their conformations appear to be much more uncorrelated than ours (see fig. 3 in ref. 24). We may attribute this difference to the fact that, in our case, the Ewald summation method is more constraining for the protein structure. Thus explaining the attraction exerted by the initial structure during the early part of the simulation.

Cluster **B** has a number of well separated subclusters showing clear PESM activity. They correspond more to the expected behaviour of a protein with the trajectory jumping between metastable substates[17,20]. Visual inspection of the distance matrix plot in fig. 4 shows that there is a clear separation between clusters **A** and **B**. The same can be said after analysing their respective tuple spectrum (figs. 3c and 3e) showing a great number of specific **b** tuples, which indicates that they lie in well separated regions of conformational space.

The dendograms built using a Hamming distance matrix reflect the degree of separation in conformational space between different trajectory segments: they are balanced binary trees over the trajectory.

One conclusion that can be drawn from the comparison between Hamming and RMSD clustering is that both measures are not equivalent: their behaviour can be radically different under a given clustering algorithm.

Another conclusion is that in the scope of the cases tested to this day, the Hamming distance measure appears to be more robust than RMSD: it can perform fairly well with smaller data sets, and gives consistent results at least with the two algorithms employed here. In the present work the RMSD seems to work well only with averaged distances, something that was also reported in ref. 7. We see three reasons for that.

First, hierarchical clustering algorithms rely upon iterative procedures, which in many cases are known to be extremely sensitive to initial conditions [25], thus the observed instability of the double linkage algorithm with the RMSD is not an uncommon behaviour.

Second, the r.m.s. deviation is built on N terms, while Hamming distance is evaluated over $\binom{N}{4}$ binary descriptors, roughly $O(N^4)$. The fact that increasing the number of atoms improves the results of RMSD calculations seems to support this conclusion. Also, the Hamming distance matrix in fig. 4 appears to be more detailed at the fine grained level than the r.m.s. deviation matrix.

Third, the hierarchical structure of conformational spaces is built-in within the Hamming distance. This follows from its physical interpretation as the minimum number of topological transitions (i.e. a point crossing the plane defined by three other), that have to be effected to transform one conformation into another. This has the effect that no matter how much we change this discrete distance by adding extra points, it remains that the old structure is always embedded within the new one. This feature helps to understand the observed good scaling behaviour.

Conclusion

In this paper we aimed at describing a new distance measure for molecular conformational space, and we used it for clustering the conformations sampled in the MD simulation of the HIV-1 integrase catalytic core.

We have described how molecular conformations can be characterized by a binary sequence: the chirotope, in which each bit is the value of a binary topological descriptor. The distance between any two conformations is the number of different bits between the corresponding sequences; this is the definition of a Hamming distance, which is a discrete one.

We have shown that, in the cases studied so far, this measure works differently than the r.m.s. deviation, which is a currently employed similarity measure, and most of the time works better. We see two reasons for this:

1) the much greater number of descriptors used in computing our distance, thus embodying far more information,

2) the fact that our topological descriptors are related to well defined geometrical structures in conformational space.

There lies the strength of the Hamming measure: it is built on meaningful structures in higher dimensional space. This is useful for analyzing the geometrical features of the clusters obtained: with the help of simple statistical data we can extract a lot of information from each one.

It appears that there are two qualitatively different regions in our trajectory. The first one is dominated by an undifferentiated cluster of conformations evolving from the initial conditions, ponctuated by incursions into substates that may be transient inclusions or appendages. In the second region the trajectory visits in succesion a series of well defined conformational substates.

This analysis seems to indicate that the equilibration time of the system may be longer than usually assumed. This same conclusion was reached in refs. 1 and 25. We suggest that the information gathered be put to use to design better MD simulation protocols.

References

1. Gabarro-Arpa, J., Revilla, R. Clustering of a molecular dynamics trajectory with a Hamming distance. Comput. and Chem. 24:693-698, 2000.

2. Noguti, T. and Gō, N. Structural basis of hierarchical multiple substates of a protein. Proteins: Struct. Funct. Genet. 5:97,104,113,125,132, 1989.

3. Frauenfelder, H., Sligar, S.G., Wolynes, P.G. The energy landscapes and motions of proteins. Science 254:1598-1603, 1991.

4. Karpen, M.E., Tobias, D.J. and Brooks III, C.L. Statistical clustering techniques for the analysis of long molecular dynamics trajectories: analysis of 2.2-ns trajectories of YPGDV. Biochemistry 32:412-420, 1993.

5. Shenkin, P.S. and McDonald, Q. Cluster analysis of molecular conformations. J. Comput. Chem. 15:889-916, 1994.

6. Torda, A.E. and van Gunsteren, W.F. Algorithms for clustering molecular dynamics configurations. J. of Comput. Chem. 15:1331-1340, 1994.

7. Troyer, J.M. and Cohen, F.E. Protein conformational landscapes: energy minimization and clustering of a long molecular dynamics trajectory. Proteins: Struct. Funct. Genet. 23:97-110, 1995.

8. Maignan, S., Guilloteau, J. P., Zhou-Liu, Q., Clément-Mella, C., Mikol, V. Crystal structures of the catalytic domain of HIV-1 integrase free and complexed with its metal cofactor: high level of similarity of the active site with other viral integrases. J. Mol. Biol. 282:359-368, 1998.

9. Pothier, J., Gabarro-Arpa, J., Le Bret, M. MORMIN: a quasi-newtonian energy minimizer fitting the nuclear Overhauser data. J. Comp. Chem. 14:226-236, 1993.

10. Pearlman, D.A., Case, D.A., Caldwell, J.C., Ross, W.S., Cheatham III, T.E., Ferguson, D.M., Seibel, G.L., Singh, U.C., Weiner, P., Kollman, P.A. AMBER 4.1, University of California, San Francisco, 1995.

11. Cornell, W. D., Cieplak, P., Bayly, C. I., Goulg, I. R., Merz, K. M., Ferguson, D. M., Spellmeyer, D. C., Fox, T., Caldwell, J. W., Kollman, P. A. A second generation force field for the simulation of proteins, nucleic acids, and organic molecules. J. Am. Chem. Soc. 117:5179-5197, 1998.

12. Aqvist, J. Ion-water interaction potentials derived from free energy perturbation simulations. J. Phys. Chem. 94:8021-8024, 1990.

13. Durden, T.A., York, D., Pedersen, L. Particle mesh Ewald: an $Nlog(N)$ method for Ewald sums in large systems. J. Chem. Phys. 98:10089, 1993.

14. Ryckaert, J.P., Cicotti, G., Berendsen, H. Numerical integration of the cartesian equations of motion of a system with constaints: molecular dynamics of n-alkanes. J. Comput. Phys. 23:327-341, 1977.

15. Ouali, M., Gousset, H., Geinguenaud, F., Liquier, J., Gabarro-Arpa, J., Le Bret, M. and Taillandier, E., Hydration of the $dT_n.dA_n \times dT_n$ parallel triple helix: a Fourier transform infrared and gravimetric study correlated with molecular dynamics simulations. Nucleic Acids Reseach 25:4816-4824, 1997.

16. Gabarro-Arpa, J., Le Bret, M., Marcouyoux, A. Cadira, an object-oriented platform for modelling molecules and analyzing simulations. Comput. and Chem. 21:343-345, 1997.

17. Becker, O.M. and Karplus, M. The topology of multidimensional potential energy surfaces: theory and application to peptide structure and kinetics. J. Chem. Phys. 106:1495-1517, 1997.

18. Nienhaus, G.U., Müller, J.D., McMahon, B.H., Frauenfelder, H., Exploring the conformational energy landscape of proteins. Physica D 107:297-311, 1997.

19. Onuchic, J.N., Luthey-Schulten, Z. and Wolynes, P.G. Theory of protein folding: the energy landscape perspective. Annual Review of Physical Chemistry 48:545-600, 1997.

20. Kitao, A., Hayward, S., Gō, N. Energy landscape of a native protein: jumping-among-minima model. Proteins: Struct. Funct. Genet. 33:496-517, 1998.

21. Krznaric, D. and Levcopoulos, C. The first subquadratic algorithm for complete linkage clustering. Proc. 6th Annu. Internat. Sympos. Algorithms Comput., Vol. 1004 of Lecture Notes in Computer Science, New-York: Springer-Verlag. 1995:392-401.

22. Kabsch, W. A discussion of the solution for the best rotation to relate two sets of vectors. Acta Cryst. A34:827-828, 1978.

23. Wlodawer, A., Deisenhofer, J., Huber, R. Comparison of two highly refined structures of bovine pancreatic trypsin inhibitor. J. Mol .Biol. 193:145-156, 1987.

24. Eastman, P., Pellegrini, M., Doniach, S. Protein flexibility in solution and in crystals. J. Chem. Phys. 110:10141-10152, 1999.

25. Ott, E., Chaos in dynamical systems. Cambridge University Press, Cambridge (England). 1993.

Cluster	Start	Length	D	$< \Delta rms >$	RMSDS
A	0	2.61	6.86		-
$\mathbf{A}_0$	0	1.4	1.94		-
$\mathbf{A}_{00}$	0	0.045	11.07	0.027 ± 0.010	yes
$\mathbf{A}_{01}$	0	0.255	7.76	0.028 ± 0.012	yes
$\mathbf{A}_1$	0.665	0.135	2.06	0.031 ± 0.010	yes
$\mathbf{A}_2$	0.79	0.16	0.88	0.034 ± 0.007	yes
$\mathbf{A}_3$	1.49	0.32	2.30	0.029 ± 0.011	yes
$\mathbf{A}_4$	2.015	0.43	0.94	0.025 ± 0.011	yes
$\mathbf{A}_5$	2.34	0.165	1.86	0.034 ± 0.009	yes
B	2.61	1.395	12.83		-
$\mathbf{B}_1$	2.61	0.22	1.49	0.027 ± 0.012	yes
$\mathbf{B}_2$	2.745	0.165	2.05	0.030 ± 0.010	yes
$\mathbf{B}_3$	2.97	0.545	0.92	0.026 ± 0.350	no
$\mathbf{B}_4$	3.54	0.29	5.33	0.020 ± 0.009	yes
$\mathbf{B}_5$	3.83	0.175	11.67	0.028 ± 0.010	yes

Table I.

In this table we show some parameters for a group of selected clusters obtained with the complete-link algoritm and the Hamming distance.

In the **first** column is the name of the cluster. In the present nomenclature subclusters inherit the name of the parent cluster, suffixed with a number corresponding to their chronological ordering inside the parent cluster.

The **second** column is the start of the cluster, in ns from the beginning of the simulation. The **third** column is the total length (in ns). Followed by the specific activity coefficient D.

The next two columns refer to the r.m.s. deviation matrix shrinkage test (RMSDS) upon minimization: we calculate two r.m.s. deviation matrices for the frames belonging to the cluster. The first one, for the original structures, the second one, for the minimized structures; in the **fifth** column there is the mean value of the difference between the elements of the original and the minimized deviation matrices.

The **last** column states the result of the RMSDS test.

Legends of Figures

Figure 1

Schematic representation of a protein conformational space. The continuous thin lines are the contour lines of two PESMs: **R1** and **R2**. The straight thick lines represent tuples that cut conformational space into positive and negative halves. The broken dotted line represents the protein trajectory, pushing the analogy a little further one can think of one frame per dot. The tuples label begins with :

b if they fall approximatively in between the two PESM,

ns if they are non specific,

s if they are specific,

sm if they cut a PESM approximatively across its center,

sw if they scratch the wall.

The approximate sequences associated with the example tuples of the figure are

b1 11111111111111111111111111110010000000000000000000000000000

b2 00000000000000000000000000000000000000010011111111111111111

ns 00000110111111111011111110111100100101000000000000001001101

sm1 10011111111111010000000000000000000000000000000000000000000

sm2 00000000000000000000000000000000000000000000000001110110000

sw1 00000001000000000000000000000000000000000000000000000000000

sw2 00000000000011000000000000000000000000000000000000000000000

sw3 00000000000000000000000000000000000000000000000010000000000

The sequences have been scaled so that each bit corresponds approximatively to 17 frames (dots) of the trajectory. For each tuple, the 1s are arbitrarily assigned to the halfspace where the trajectory spends the less time.

Figure 2

Dendogram resulting from the complete-link clustering algorithm with the Hamming distance. It has been drawn so that the ordering of the clusters along the horizontal axis remains close to the chronological ordering. The clusters found with our analysis protocol are underlined. The names of subclusters begin with parent cluster's name followed by a number corresponding to their time ordering within the cluster. For sake of clarity the diagram has been cut for an arbitrary distance since the lower branches do not convey important information.

Figure 3

Spectrum of specific 4-tuple for some selected clusters.

The vertical coordinate, in logarithmic scale, is the total number of tuples. The horizontal coordinate is the trajectory total number of frames in the tuple minor halfspace.

For each cluster two curves are plotted :

- the full line is the spectrum of the specific tuples;
- the dotted line is the average number spectrum of specific tuples on the homologous clusters along the trajectory.

Figure 4

Matrix density plot.

Lower left half: r.m.s. deviation matrix. Upper right half: Hamming distance matrix.

Axis coordinates are all in ns.

Figure 5

a) Dendogram resulting from the complete-link clustering algorithm with the r.m.s. distance.

b) Dendogram resulting from the average distance clustering algorithm with the r.m.s. distance.

For comparison we have underlined the clusters corresponding to clusters or segments of clusters identified from the dendogram obtained with the Hamming distance. The symbols have the same meaning as in table I and fig. 2.For sake of clarity the diagram has been cut for an arbitrary distance since the lower branches do not convey important information.

Clusters bearing a $'$, have more than 80% homology with their homologous counterparts in fig. 2.

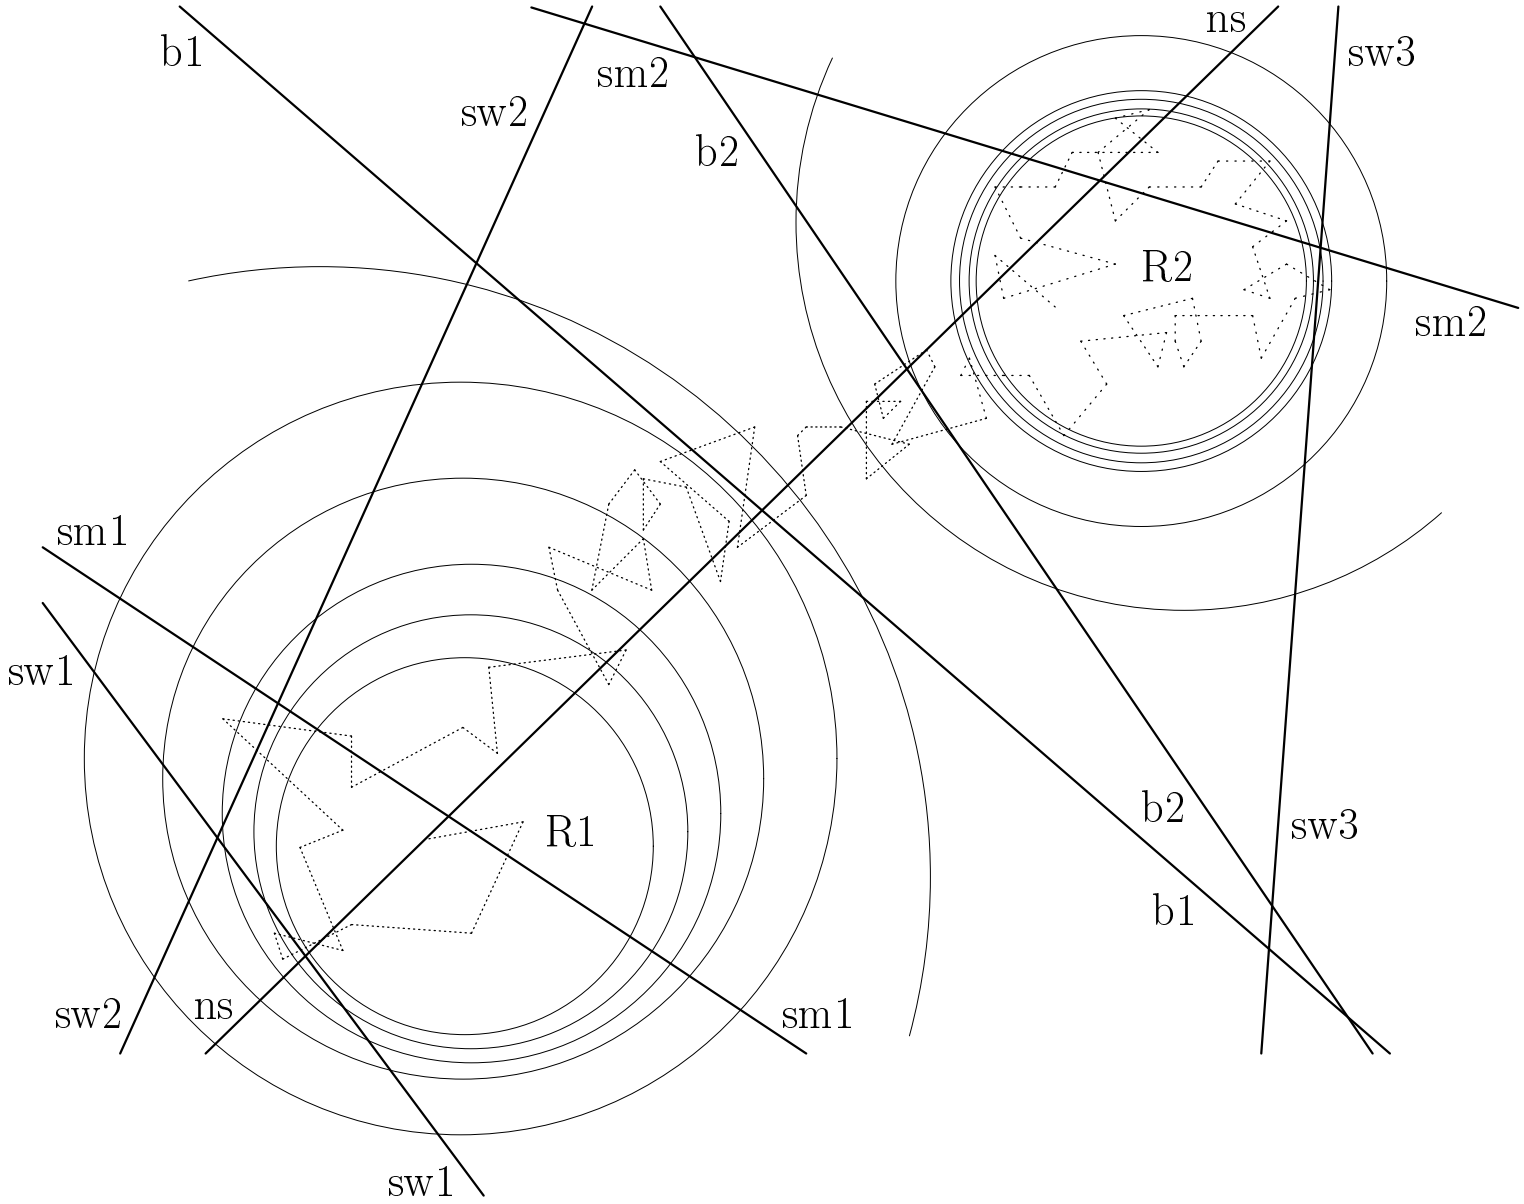

Figure 1

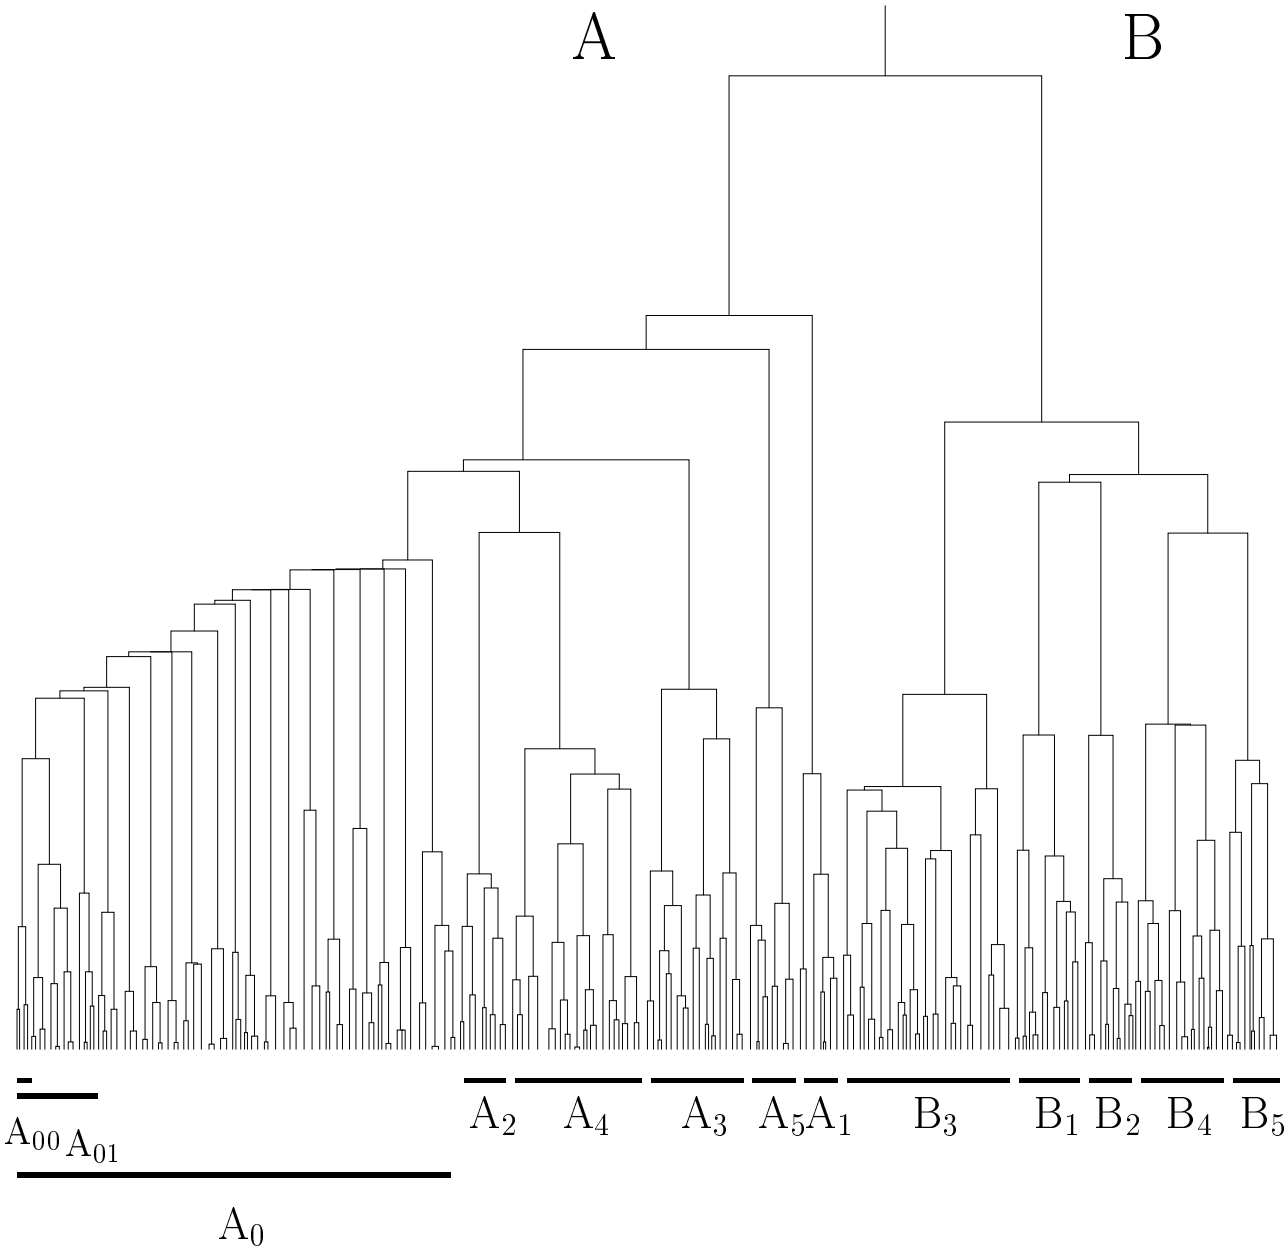

Figure 2

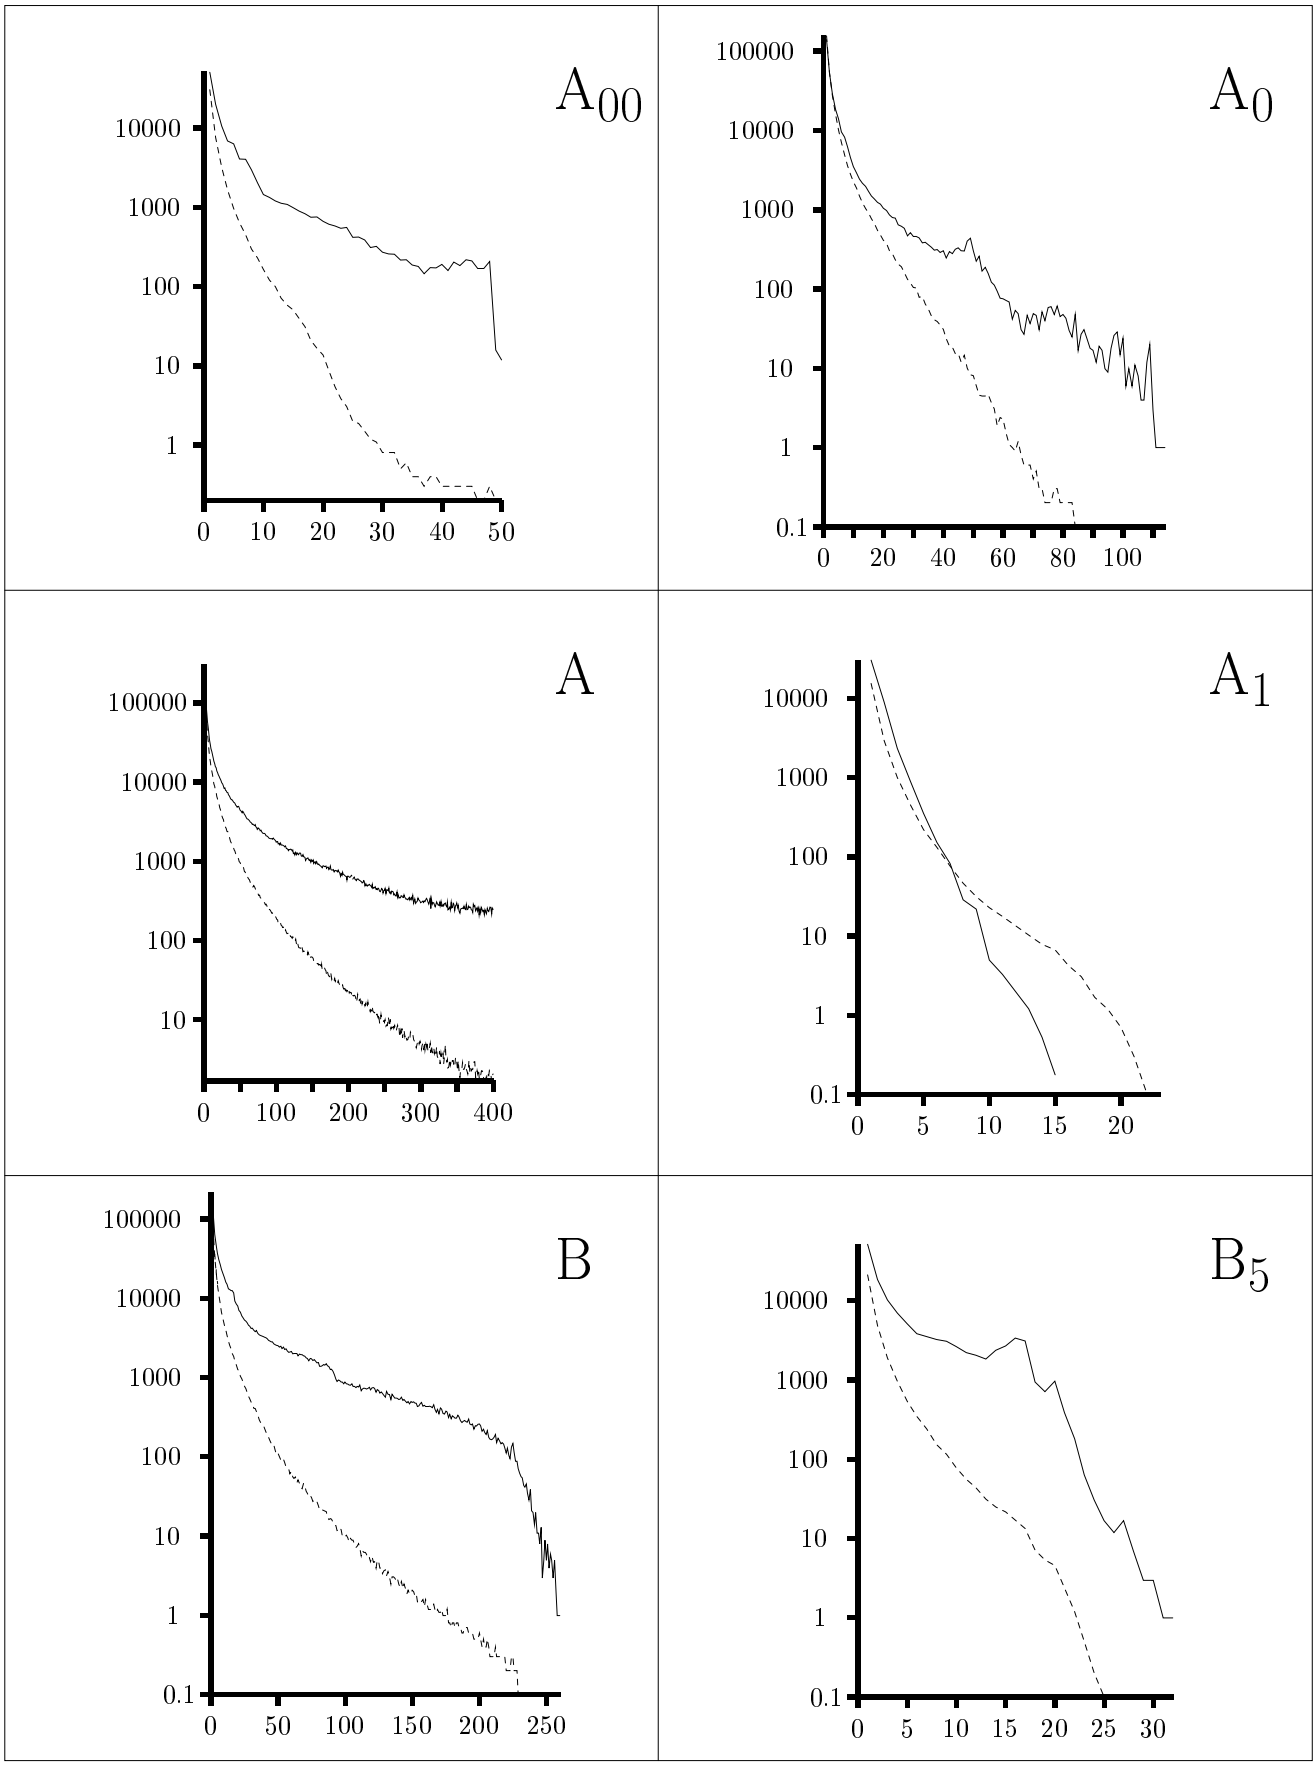
A00
10000
1000
100
10
1
0 10 20 30 40 50
A0
100000
10000
1000
100
10
1
0.1
0 20 40 60 80 100
A
100000
10000
1000
100
10
0 100 200 300 400
A1
10000
1000
100
10
1
0.1
0 5 10 15 20
B
100000
10000
1000
100
10
1
0.1
0 50 100 150 200 250
B5
10000
1000
100
10
1
0.1
0 5 10 15 20 25 30

Figure 3

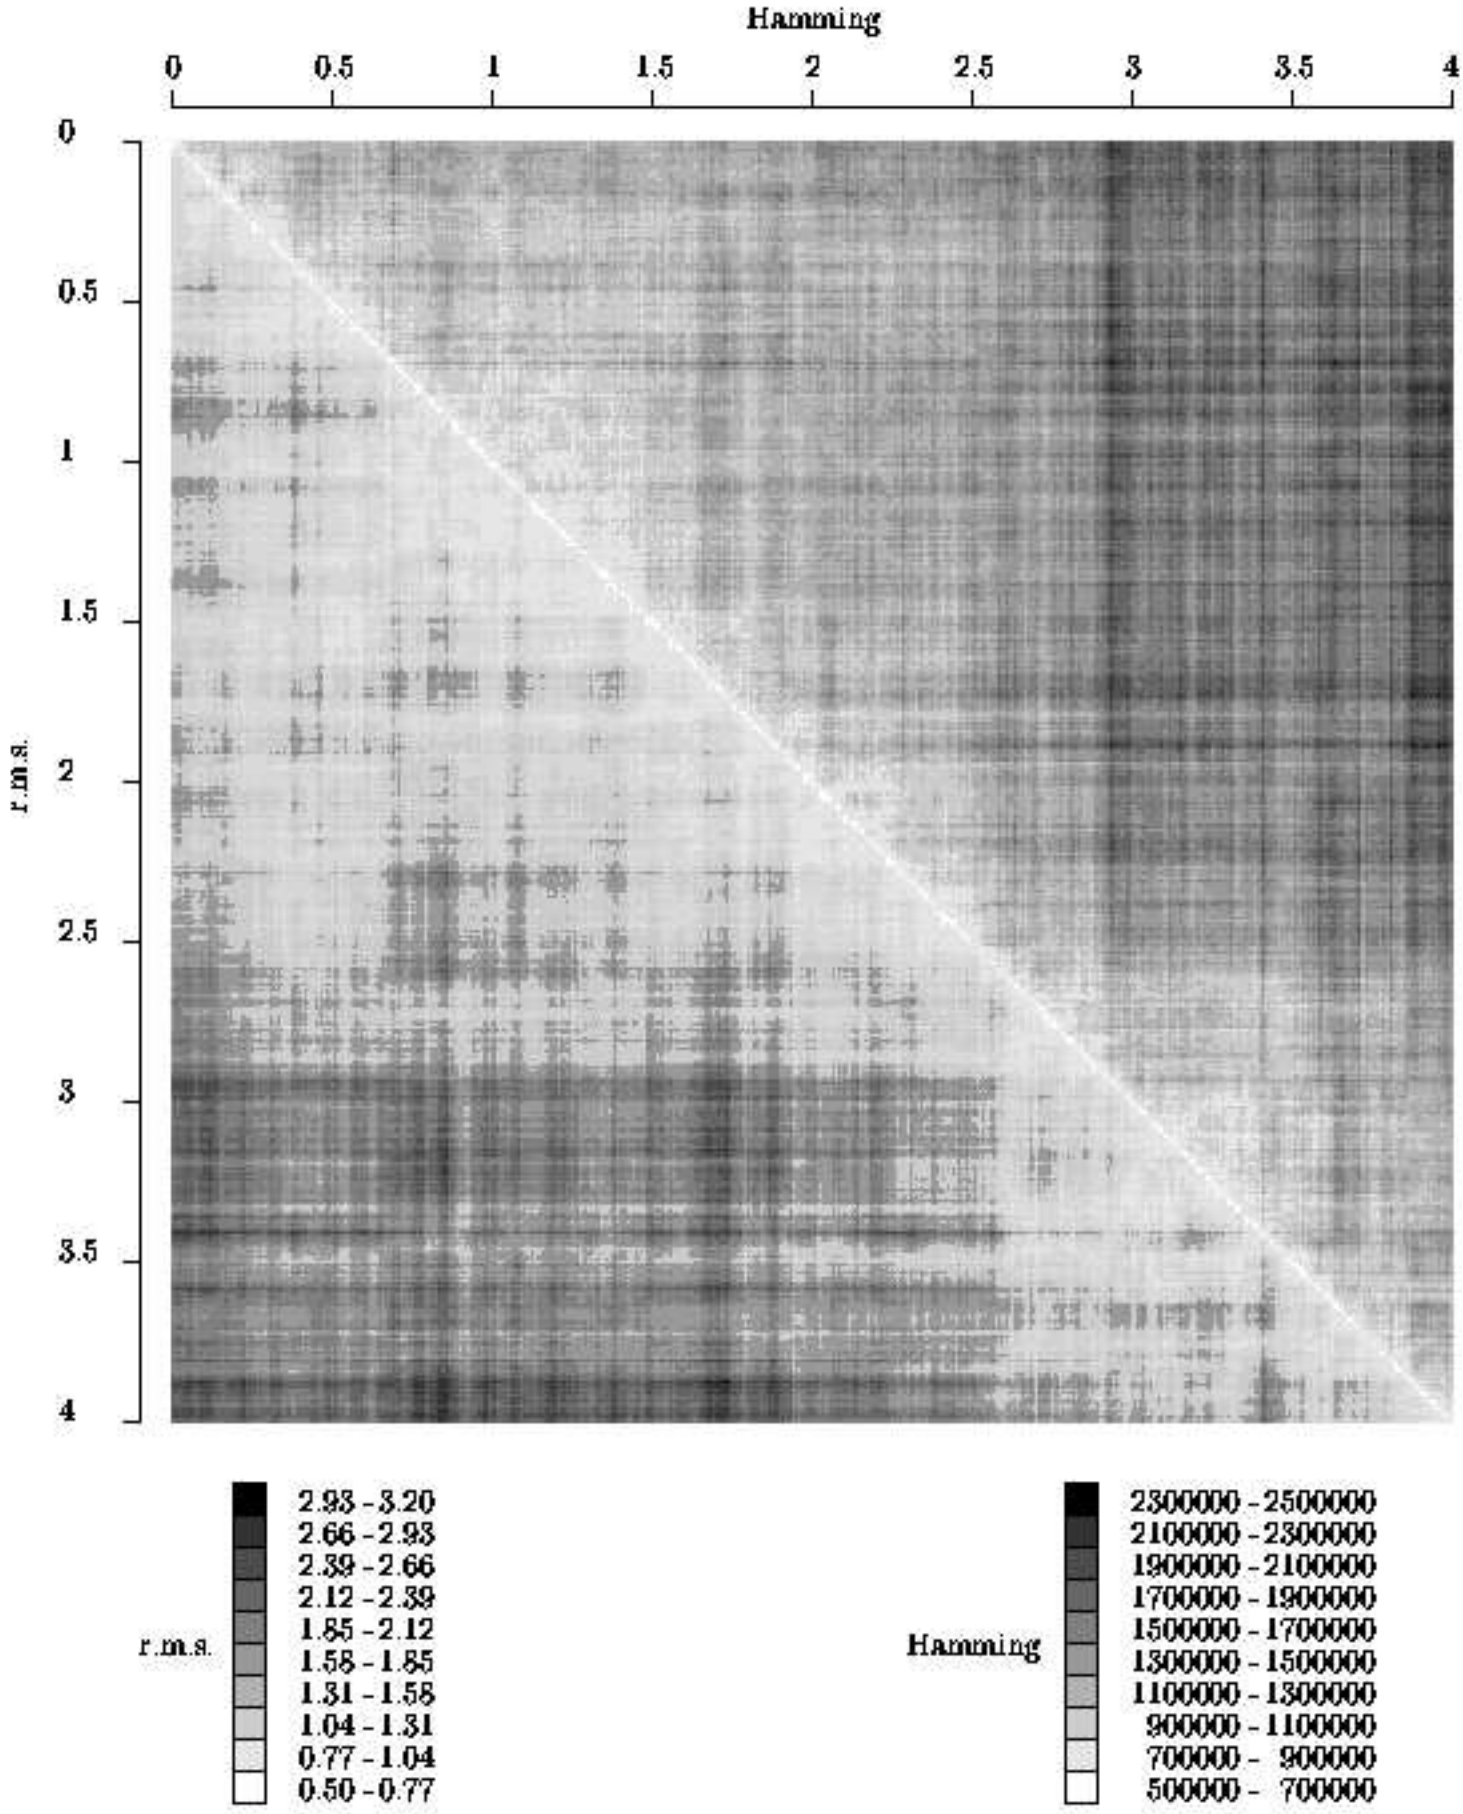
Hamming
0
0.5
1
1.5
2
2.5
3
3.5
4
r.m.s.
r.m.s.
2.93 - 3.20
2.66 - 2.93
2.39 - 2.66
2.12 - 2.39
1.85 - 2.12
1.58 - 1.85
1.31 - 1.58
1.04 - 1.31
0.77 - 1.04
0.50 - 0.77
Hamming
2300000 - 2500000
2100000 - 2300000
1900000 - 2100000
1700000 - 1900000
1500000 - 1700000
1300000 - 1500000
1100000 - 1300000
900000 - 1100000
700000 - 900000
500000 - 700000

Figure 4

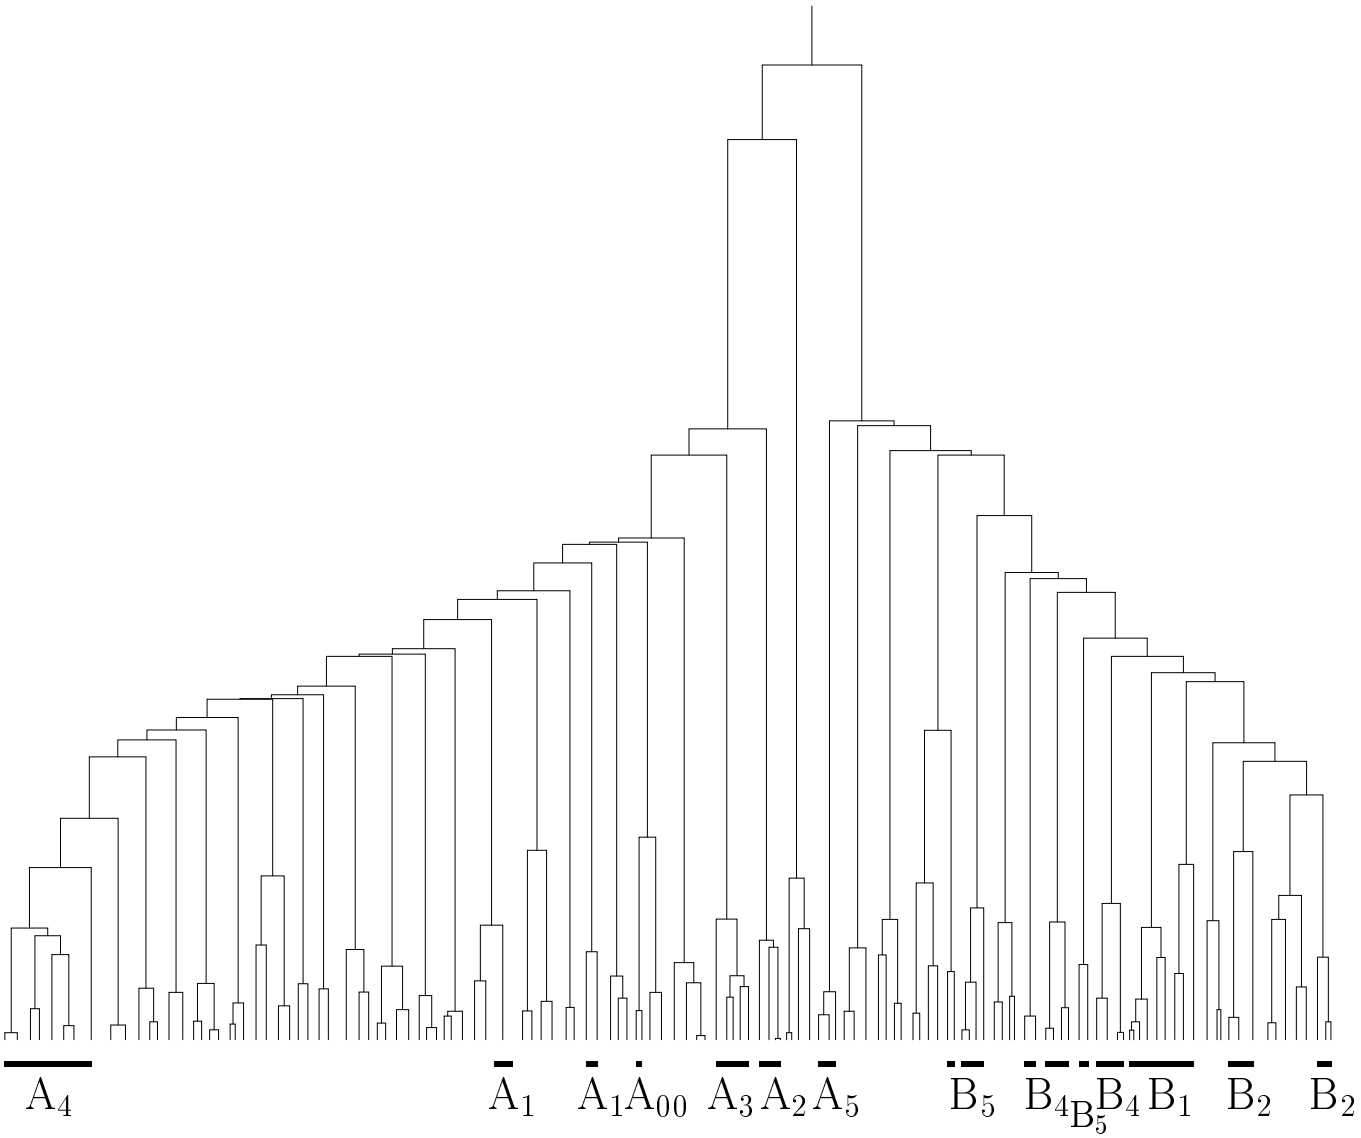

Figure 5a

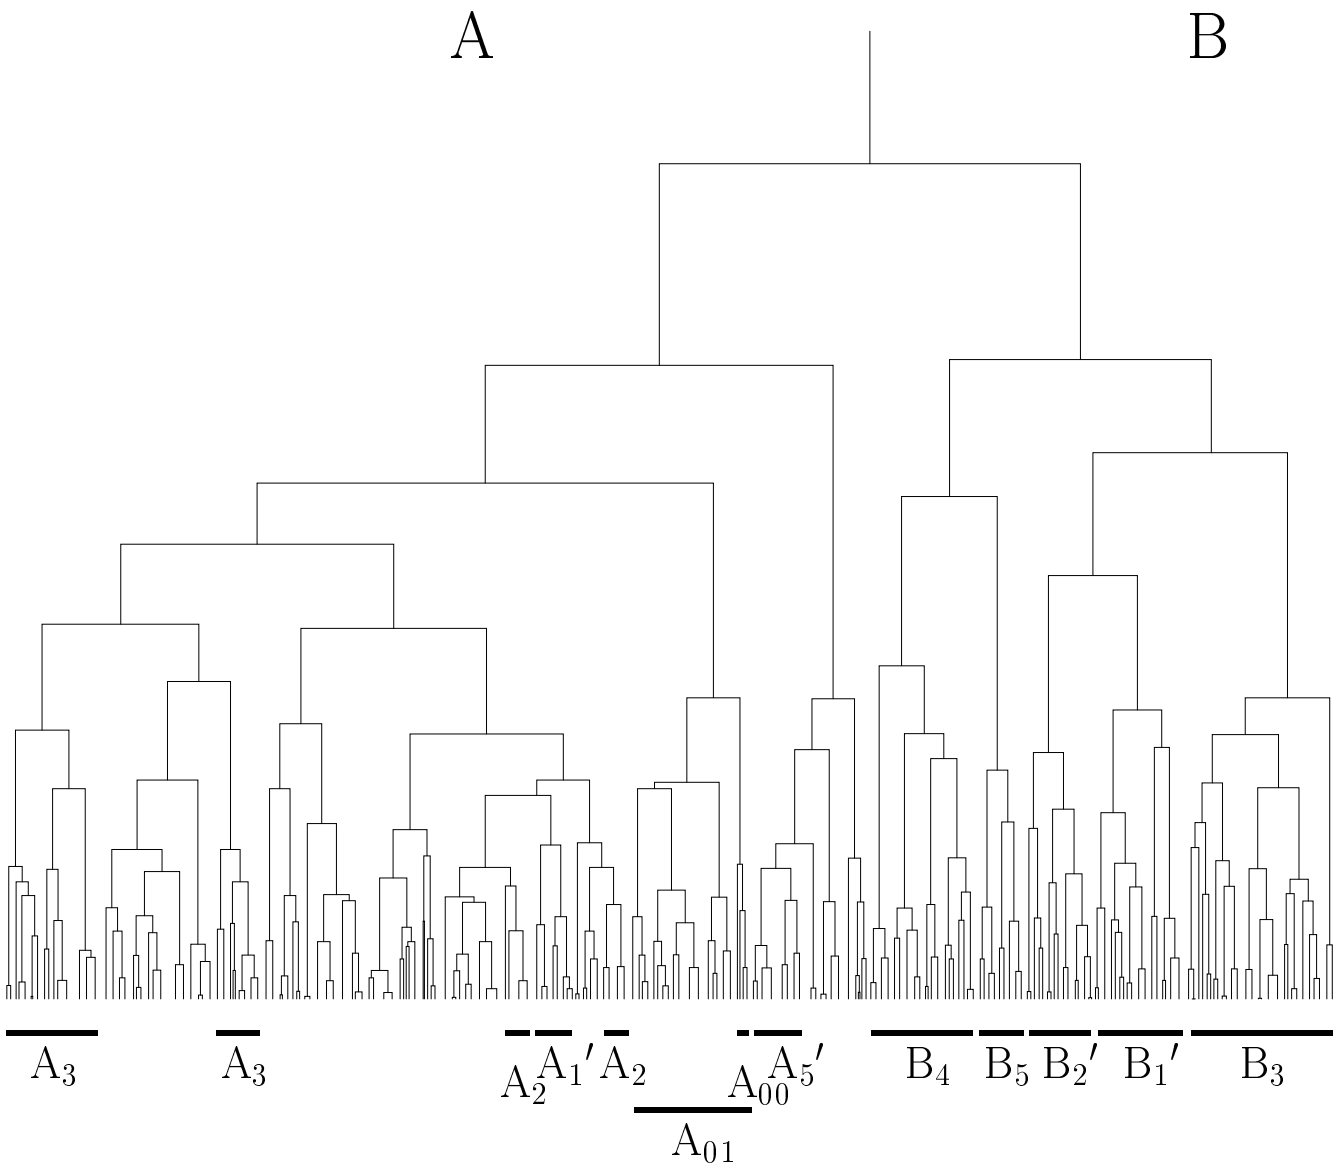

Figure 5b